

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
19 September 2002 (19.09.2002)

PCT

(10) International Publication Number
WO 02/073440 A1

(51) International Patent Classification⁷: G06F 15/16

(21) International Application Number: PCT/US02/07659

(22) International Filing Date: 12 March 2002 (12.03.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/275,407 12 March 2001 (12.03.2001) US
60/275,408 12 March 2001 (12.03.2001) US

(71) Applicant (for all designated States except US):
EDGESTREAM, INC. [US/US]; Suite 231, 22931
Triton Way, Laguna Hills, CA 92653 (US).

(72) Inventors; and

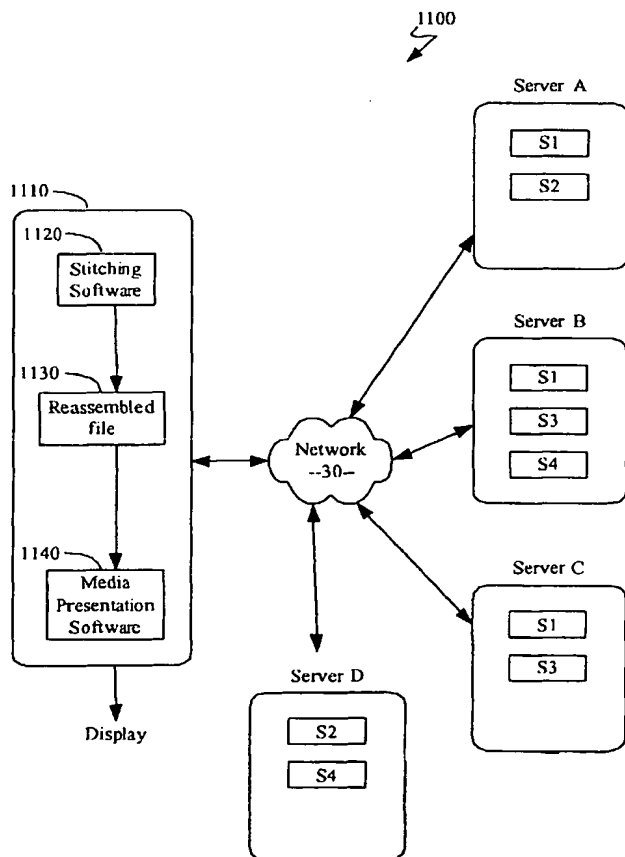
(75) Inventors/Applicants (for US only): CHUNG, Randall, M. [US/US]; 30542 Via Estoril, Laguna Niguel, CA 92677 (US). BISHAY, Maged [US/US]; 3223 Colorado Place, Costa Mesa, CA 92626 (US). ZELEVINSKY, Michael [US/US]; Apartment 48, 7164 Schilling Avenue, San Diego, CA 92126 (US). PARTHASARATHY, Sivagnanam [US/US]; 2187 Corte Mango, Carlsbad, CA 92009 (US).

(74) Agent: REDJAIAN, Babak; Irell & Manella LLP, Suite 400, 840 Newport Center Drive, Newport Beach, CA 92660 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX,

[Continued on next page]

(54) Title: RE-ASSEMBLY OF STREAMING FILES FROM SEPARATE CONNECTIONS



(57) Abstract: The invention is a technique and apparatus for re-assembling transmitted portions of files or sub-files (410, 420, 430, 440), into a viewable file (400). The transmitting device passes the information regarding the file dividing process to the receiver so that the sub-files can be assembled (1130) or stitched (1120) together properly. The media player (1140) in the receiver uses the attached instructions to assemble the sub-files into one whole piece for viewing.

WO 02/073440 A1



MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL,
TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for all designations*
- *of inventorship (Rule 4.17(iv)) for US only*

Published:

- *with international search report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

RE-ASSEMBLY OF STREAMING FILES FROM SEPARATE CONNECTIONS

BACKGROUND OF THE INVENTION

The present invention claims priority to U.S. provisional patent application serial number 60/275,407, entitled "Re-Assembly of Streaming Files from Separate Connections," filed March 12, 2001, and U.S. provisional patent application serial number 60/275,408, entitled "Splitting and Redundant Storage on Multiple Servers", filed March 12, 2001. The present application is also related to pending PCT Application No. PCT/US___/_____, entitled "Splitting and Redundant Storage on Multiple Servers," filed March 12, 2002, which is assigned to the assignee of the present application, the subject matter of which is incorporated herein by reference.

1. Field of the Invention

The present invention relates generally to data communication networks and more particularly to a method and apparatus for reassembling streaming files from separate connections.

2. Related Art

Replicating an entire file onto multiple file servers is a technique that can be used to speed up the delivery of data. Once the data has been replicated onto multiple servers, and the servers have been geographically dispersed, the data can be delivered to an end user from the closest available server. In addition, the data can be delivered from multiple servers simultaneously, where each server provides a portion of the

total file. However, replicating the full file onto a large number of servers uses large amounts of expensive disk storage. For instance, a system consisting of a hundred servers with 100 gigabytes of storage each would have a total of 10,000 gigabytes of storage available if data were not replicated. If all files must be fully replicated onto each server, however, the system would only be able to store 100 gigabytes of unique files.

U.S. provisional patent application serial number 60/275,408, entitled "Splitting and Redundant Storage on Multiple Servers", filed March 12, 2001, describes a method and apparatus for dividing a file into multiple sub-files and storing the sub-files onto one or more servers. The servers may be distributed over different geographical locations. When an end user wants the file to be delivered in a streaming fashion, the sub-files can be transmitted in parallel from one or more servers, which increases the rate at which data can be delivered.

While splitting and distributed storage provides many advantages, it creates a daunting re-assembly task for the end-user's machine.

Accordingly, there is a need in the technology for a method and apparatus for overcoming the aforementioned problems.

BRIEF SUMMARY OF THE INVENTION

The invention is a technique and apparatus for re-assembling transmitted portions of files or sub-files, into a viewable file. The transmitting device passes the information regarding the file dividing process to the receiver so that the sub-files can be assembled or stitched back together properly. The media player in the receiver

uses the attached instructions to assemble the sub-files into one whole piece for viewing.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a system block diagram of one embodiment of a network system in which the system and method of the invention may be implemented.

Figure 2 is a system block diagram of one embodiment of a computer system which implements the embodiments of the invention.

Figure 3 is a flow chart illustrating one embodiment of the setup process provided in accordance with the principles of the invention.

Figure 4 illustrates one embodiment of a process for splitting and distributing a file, in accordance with the principles of the invention.

Figure 5 illustrates one embodiment of a process for replicating storage of sub files, provided in accordance with the principles of the invention.

Figure 6 is a flow chart illustrating one embodiment of a process flow for configuring a file for storage, provided in accordance with the principles of the invention.

Figure 7 illustrates one embodiment of the beginning portion of a playback process of a re-assembled media file.

Figure 8 illustrates one embodiment of a simultaneous transfer process, provided in accordance with the principles of the invention.

Figure 9 illustrates one embodiment of a transfer flow from a server, provided in accordance with the principles of the invention.

Figure 10A illustrates one embodiment of a list of server addresses, provided in accordance with the principles of the invention.

Figure 10B illustrates one embodiment of an initial server assignment, provided in accordance with the principles of the invention.

Figure 10C illustrates one embodiment of server re-assignment, provided in accordance with the principles of the invention.

Figure 11 is a system block diagram of one example of a networked system that implements one embodiment of the process of the invention.

DETAILED DESCRIPTION OF THE EMBODIMENTS

U.S. provisional patent application serial number 60/275,408, entitled "Splitting and Redundant Storage on Multiple Servers", filed March 12, 2001, the subject matter of which is incorporated herein by reference, describes the division of a single file into multiple sub-files and storage the sub-files onto one or more servers. Such a file may include text, data, graphics, video clips, JPEG elements or images, static photographs, web pages, audio clips, animation, any type of informational material or any combination thereof. The servers may be distributed over different geographical locations. When an end user wants the file to be delivered in a streaming fashion, the sub-files can be transmitted in parallel from one or more servers, which increases the rate at which data can be delivered.

In one embodiment, each sub-file may be replicated onto more than one server so as to provide some redundancy. If one server is not available, or if the transmission link is slow or not available, the sub-file can be streamed to the end user from another server.

Replicating a file onto multiple file servers in accordance with the principles of the invention facilitates the transmission rate of the data. For instance, the data may be delivered to an end user from the closest and fastest server. In addition, the data may be delivered from multiple servers simultaneously, where each server provides a portion of the total file.

One aspect of the invention involves a technique and apparatus that may be used to re-assemble the sub-files into the viewable file. Such a technique may be implemented, for example, by the receiver or by the end-user's machine. The server passes the information regarding the splitting process to the receiver or end-user device so that the sub-files can be assembled or stitched back together properly. The media player in the receiver will use the attached instructions to stitch the sub-files into one whole piece for viewing.

A second aspect of the invention involves the determination of the optimum servers to be used for the transmission, if more than one server can actually provide the requested sub-file, and to change to the optimum servers at any time.

Definitions

As discussed herein, a "computer system" is a product including circuitry capable of processing data. The computer system may include, but is not limited to,

general purpose computer systems (e.g., server, laptop, desktop, palmtop, personal electronic devices, etc.), personal computers (PCs), hard copy equipment (e.g., printer, plotter, fax machine, etc.), banking equipment (e.g., an automated teller machine), and the like. Content and/or media elements refers to application programs, driver programs, utility programs, file, payload, and combinations thereof, as well as graphics, informational material (articles, stock quotes, etc.) and the like, either singly or in any combination. A "communication link" refers to the medium or channel of communication. The communication link may include, but is not limited to, a telephone line, a modem connection, an Internet connection, an Integrated Services Digital Network ("ISDN") connection, an Asynchronous Transfer Mode (ATM) connection, a frame relay connection, an Ethernet connection, a coaxial connection, a fiber optic connection, satellite connections (e.g. Digital Satellite Services, etc.), wireless connections, radio frequency (RF) links, electromagnetic links, two way paging connections, etc., and combinations thereof.

System Overview

A description of an exemplary system, which incorporates embodiments of the invention, is herein described. Figure 1 shows a system block diagram of one embodiment of a network system 10 in which the apparatus and method of the invention is used. Referring to Figure 1, the network system 10 comprises a service center 12 that is connected over one or more communication links 20 to a remote network 30 (e.g., a wide area network or the Internet) or a remote site (e.g., a satellite, which is not shown in Figure 1) to one or more user computer systems 40₁-40_N ("40"). In one embodiment the service center 12 is a website. The service center 12

includes one or more servers 22 and one or more databases 24. In one embodiment, the server 22 includes software modules that may be downloaded for performing the processes of the invention, as described in detail in the following sections. In addition, there may be more than one service center 12, each having one or more servers.

The service center 12 may also include one or more computers 26₁-26_M. If a plurality of computers are used, then the computers 26₁-26_M may be connected by a local area network (LAN) or any other similar connection technology. However, it is also possible for the service center 12 to have other configurations. For example, a smaller number of larger computers (i.e. a few mainframe, mini, etc. computers) with a number of internal programs or processes running on the larger computers capable of establishing communication links to the user computers.

The remote network 30 or remote site allows the service center 12 to provide peer-to-peer information and services to the user computers 40₁-40_N, using software that is stored at the service center 12. In one embodiment, the user computers 40 may be any computer system or any other device that processes and/or presents audio and/or video information. The one or more databases 24 connected to the service center computer(s), e.g., computer 26₁, are used to store data. Each user computer 40 is connected via network connection 32₁-32_N over a corresponding communication link 42₁-42_N such as a local carrier exchange to a respective ISP 44₁-44_N, through which access to the remote network 30 is made.

By inputting the URL address of the target website with which the user desires to interact, the user may be connected to various target websites, such as

websites 50₁-50_{NN}. In an alternate embodiment, each user computer 40 may be connected via network connection 32₁-32_N over a corresponding communication link 48₁-48_N to the service center 12, which provides internet access and service to the user computer(s) 40. In a further embodiment, the display screen for viewing the content or media elements may be located on a television coupled to the network 30. In this case, navigation through the content or media elements may be provided through the use of control buttons on a remote control unit for controlling viewing of the television, or by other means known in the art.

One aspect of the present invention relates to retrieval and delivery of content and/or media elements. The software for providing such processes may be developed and/or stored on a computer 40, or may be developed using one of the computers 26₁-26_M. Upon completion of the development process, the software may be stored in the computer 40 or in the database 24. Alternatively, the software may be stored on a machine-readable medium, in any of the computers 40 or computers 26₁-26_M.

Referring to Figure 2, the computer system 100 (representing either of computer 26 or 40) comprises a processor or a central processing unit (CPU) 110. The illustrated CPU 110 includes an Arithmetic Logic Unit (ALU) for performing computations, a collection of registers for temporary storage of data and instructions, and a control unit for controlling operation for the system 100. In one embodiment, the CPU 110 includes any one of the x86, Pentium™, Pentium II™, and Pentium Pro™ microprocessors as marketed by Intel™ Corporation, the K-6 microprocessor as marketed by AMD™, or the 6x86MX microprocessor as marketed by Cyrix™ Corp. Further examples include the Alpha™ processor as marketed by Digital

Equipment Corporation™, the 680X0 processor as marketed by Motorola™; or the Power PC™ processor as marketed by IBM™. In addition, any of a variety of other processors, including those from Sun Microsystems, MIPS, IBM, Motorola, NEC, Cyrix, AMD, Nexgen and others may be used for implementing CPU 110. The CPU 110 is not limited to microprocessor but may take on other forms such as microcontrollers, digital signal processors, reduced instruction set computers (RISC), application specific integrated circuits, and the like. Although shown with one CPU 110, computer system 100 may alternatively include multiple processing units.

The CPU 110 is coupled to a bus controller 112. The bus controller 112 may include a memory controller (not shown) integrated therein, though the memory controller may be external to the bus controller 112. The memory controller provides an interface for access by the CPU 110 or other devices to memory 116 via memory bus 114. In one embodiment, the system memory 116 includes synchronous dynamic random access memory (SDRAM). System memory 116 may optionally include any additional or alternative high speed memory device or memory circuitry. The bus controller 112 is coupled to a system bus 120 that may be a peripheral component interconnect (PCI) bus, Industry Standard Architecture (ISA) bus, etc. Coupled to the system bus 120 are a graphics controller, a graphics engine or a video controller 132, a mass storage device 150, a communication interface device 152, and one or more input/output (I/O) devices 168₁-168_N. The video controller 132 is coupled to a video memory 136 (e.g., 8 Megabytes) and video BIOS 140, all of which may be integrated onto a single card or device, as designated by numeral 144. The video memory 136 is used to contain display data for displaying information on the display screen 148, and the video BIOS 140 includes code and video services for controlling the video

controller 132. In another embodiment, the video controller 132 is coupled to the CPU 110 through an Advanced Graphics Port (AGP) bus.

The mass storage device 150 includes (but is not limited to) a hard disc, floppy disc, CD-ROM, DVD-ROM, tape, high density floppy, high capacity removable media, low capacity removable media, solid state memory device, etc., and combinations thereof. The mass storage device 150 may include any other mass storage medium. The communication interface device 156 includes a network card, a modem interface, etc. for accessing network 164 via communications link 160. The I/O devices 168₁-168_N include a keyboard, mouse, audio/sound card, printer, and the like. The I/O devices 168₁-168_N may be disk drive, such as a compact disc (CD) drive, a hard disc drive, a tape drive, a zip drive, a jazz drive, a digital versatile disc (DVD) drive, a magneto-optical disk drive, a high density floppy drive, a high capacity removable media drive, a low capacity media device, and/or any combination thereof.

The information, content and/or application software may be stored in the database 24, on the computers 40, or on a machine-readable medium. The information, content and/or application software may also be made available to users such as those located at user computer 1 – N, i.e., computers 40₁-40_N, through service center 12 or by means of the machine-readable medium.

As discussed earlier, one aspect of the invention relates to a system and method for providing media over a network in a non-sequential fashion. In one embodiment, a request may be made by a user for content or media elements. The system and process of the invention may retrieve the desired content/media element

from either database 24 or one of target websites 50₁-50_{NN}. The retrieved content/media element is then processed (e.g., decompressed, formatted, etc.). Users may then view the processed content/media element on a display device, such as a user computer 40. In another embodiment, the desired content/media is provided to the user in a non-sequential fashion.

In accordance with the practices of persons skilled in the art of computer programming, the invention is described below with reference to symbolic representations of operations that are performed by a computer system or a like electronic system. Such operations are sometimes referred to as being computer-executed. It will be appreciated that operations that are symbolically represented include the manipulation by a processor, such as a central processing unit, of electrical signals representing data bits and the maintenance of data its at memory locations such as in system memory, as well as other processing of signals. The memory locations where data its are maintained are physical locations that have particular electrical, magnetic, optical, or organic properties corresponding to the data bits. Thus, the term "server" is understood to include any electronic device that contains a processor, such as a central processing unit.

When implemented in software, the elements of the invention are essentially the code segments to perform the necessary tasks. The program or code segments can be stored in a processor readable medium or transmitted by a computer data signal embodied in a carrier wave over a transmission medium or communication link. The "processor readable medium" or "machine-readable medium" may include any medium that can store or transfer information. Examples of the processor

readable medium include an electronic circuit, a semiconductor memory device, a read only memory (ROM), a flash memory, an erasable ROM (EROM), a floppy diskette, a compact disk ROM (CD-ROM), an optical disk, a hard disk, a fiber optic medium, etc. The computer data signal may include any signal that can propagate over a transmission medium such as electronic network channels, optical fibers, air, electromagnetic links, a radio-frequency link, etc. The codes segments may be downloaded via computer networks such as the Internet, Intranet, etc.

Setup Process

Referring now to Figure 3, in which a setup process is depicted. Setup process 300 begins with the development of the client-side software at block 305. As will be discussed in more detail below, the client-side software will enable a user to communicate with client-side software by issuing read requests to one or more servers, according to one embodiment. In another embodiment, the client-side software also comprises stitching software which may be used to reassemble a requested file, or portion thereof, which has been previously split. One example of a method for reassembling streaming files is provided in related PCT Application No. PCT/US___/_____, entitled "Re-assembly of Streaming Files from Separate Connections" (hereinafter the "PCT/US___/_____ patent application"), which describes the reassembly of the sub-files into a viewable file upon receipt by a client system. In one embodiment, the server passes the information regarding the splitting process to the client system to facilitate the sub-file reassembly process.

Server-side software is developed at block 310. As will be discussed in more detail below, the server-side software will be able to process random read requests,

according to one embodiment. The server-side software of block 310 may also enable a server to perform offset computations, thereby enabling the server to stream content to a client system beginning at the requested location. The server-side software may also comprise software to split a single file into multiple sub-files, which may then be distributed and stored on one or more servers. One example of a method for splitting files into sub-files is provided in related PCT Application No. PCT/US__/_____, entitled "Splitting and Redundant Storage on Multiple Servers" (hereinafter the "PCT/US__/_____ patent application"), which describes the division of a single file into multiple sub-files which are then stored on one or more servers. Such a file may include text, data, graphics, video clips, JPEG elements or images, static photographs, web pages, audio clips, animation, any type of informational material or any combination thereof. The servers may be distributed over different geographical locations. In one embodiment, the sub-files can be transmitted in parallel from one or more servers in response to a user request, thereby increasing the rate at which data can be delivered.

While in the embodiment of Figure 3 the development of the server-side software takes place at block 310, it should be appreciated that the server-side software may be developed before the client-side software of block 305. It should further be appreciated that the development process of block 305 or block 310 may take place on any one of user computers 40, computers 26₁-26_M, server 22, or any other system capable of being used for software development. It is also possible to install both client-side and server-side software on the same client computer to implement a peer-to-peer system.

Continuing to refer to Figure 3, the software of block 305 and 310 are installed at block 315. It should be appreciated that the client-side software of block 305 may be installed on one or more user computers 40, while the server-side software of block 310 may be installed on server 22. In another embodiment the server-side software is installed on a plurality of servers connected either to the network 30 or to the client systems directly over a communications link.

At block 320, media presentation software is installed on one or more client system, such as one or more user computers 40. Media presentation software may include any of Windows Media Player, Real Player, QuickTime Player, or any other media players capable of receiving streaming content. It should also be appreciated that the media presentations software installation of block 320 may occur before the installation activities of blocks 315, and even before the development activities of blocks 305 and 310. Thereafter, the client system(s) on which the client-side software was loaded may be connected to the server(s) running the server-side software. In one embodiment, the client system communicates with the server(s) by connecting to the World Wide Web, while in another embodiment the client system(s) and the server(s) communicate over a LAN connection. However, it should be appreciated that the client system(s) may interact with the server(s) over any known communication link.

Figure 4 illustrates one embodiment of a process for dividing a file, in accordance with the principles of the invention. In this embodiment, a 128 kilobyte file is divided and distributed. The size of the file is merely illustrative, and it is understood that the invention may be implemented using a file of a greater or smaller

size. Such a file may include text, data, graphics, video clips, JPEG elements or images, static photographs, web pages, audio clips, animation, any type of informational material or any combination thereof.

Prior to transmission, the original file is configured for transmission. For example, the original file may be divided into N segments or chunks, where N is a positive integer. The N segments are assembled into M sub-files, where M is a positive integer. Each of the sub-files has a file name and other file attributes, and is treated by the operating system's file system as just another file. In the present example, the original 128 kilobyte file is split up into sixteen 8 KB segments (i.e., $N = 16$ in the example). The sixteen segments are assembled into four sub-files (i.e., $M = 4$ in the example), where the first 8 KB chunk of the original file goes into the first sub-file S1, the second chunk of the original file goes into the second sub-file S2, and so on. The fifth chunk of the original file would go into the first sub-file S1, after the first chunk of the original file. Thus, in one embodiment, the segments may be assigned to sub-files on a modulo 4 arrangement. Although in the present embodiment, the file is divided into 8 kilobyte segments, it is understood that the size of the segments and the number of sub-files is arbitrary. The invention may be implemented using a larger or smaller size chunk and a larger or smaller number of sub-files. In addition, each of the sub-files may be of different sizes. However, in one embodiment, the segments may be distributed in modulo order and the sub-files may be of substantially equal size.

The four sub-files S1 through S4 may be distributed to multiple servers. As shown in the example of Figure 5, the first sub-file S1 may be distributed to three

servers, servers A, B, and C. The second sub-file S2 may be distributed to servers D and A. The third sub-file S3 may be distributed to servers B and C. The fourth sub-file S4 may be distributed to servers D and B. It is understood that the sub-files may be distributed to a fewer or greater number of servers. Each sub file may be transmitted over a separate communication link, such as a TCP/IP connection to a receiver. In one embodiment, the Hyper Text Transfer Protocol (HTTP) may be used for each connection.

In one embodiment, the sub-files can be replicated so that the most frequently accessed files and sub-files are replicated onto more servers, while the less frequently accessed files and sub-files are replicated onto fewer servers. Thus, the sub files may be replicated onto multiple servers based on anticipated or actual demand. This saves disk space, since infrequently accessed files would take up less total disk space. Disk input/output bandwidth on each server would also be saved. For example, each server would read a 32 kilobyte sub-file and transmit the sub-file to the end user. If the files were not split up beforehand, each server would read the entire 128 kilobyte file, while transmitting only 32 kilobytes.

Figure 6 is a flow chart illustrating one embodiment of a process flow for configuring a file for storage. Proceeding from a START state, the process 600 proceeds to configure an original file, such as file 400 of Figure 4, for transmission. The process 600 begins by determining the number and size of segments that the original file 400 should be divided into, as shown in process block 620. The process 600 then determines the size and number of sub-files that the segments should be

assembled into (process block 630). Next, the process 600 determines the order of assembling the segments into each sub-file (process block 640).

The process 600 is then ready to divide the original file, and proceeds to configure the original file by dividing it into segments, and thereafter assembling the segments into the sub-files, as discussed above (process block 650). The process 600 then determines how the sub-files should be distributed to various servers for storage (process block 660). When this has been determined, the process 600 distributes and stores the sub-files onto the servers (process block 670). The process 600 then terminates or returns to the main process.

Figure 7 illustrates one embodiment of the beginning portion of a playback process of a re-assembled media file. For illustrative purposes, the playback process is described with reference to an end user's personal computer (PC). It is understood that the playback process and other processes described herein may be implemented on the receiver device in other electronic systems. As shown in Figure 7, the end user would view a web page that contains a link to a video clip. The link may be of the form "etp://1.2.3.4/fileS.cfg". When the end user clicks on the web page link, the web browser receives the file named file1.cfg from the server with the Internet address of 1.2.3.4, which will start the stitching program associated with the "etp:" protocol, and which will then transfer the contents of fileS.cfg to the stitching program.

Prior to transmission, the original file is configured for transmission. For example, the original file may be divided into N segments, where N is a positive integer. The N segments are assembled into M sub-files, where M is a positive

integer. In one embodiment, the first chunk of the original file goes into the first sub-file S1, the second chunk of the original file goes into the second sub-file S2, and so on. The fifth chunk of the original file would go into the first sub-file S1, after the first chunk of the original file.

The contents of fileS.cfg may contain information such as the size of the segments, the number N of the segments, the number M of sub-files, the file size of the original file, and a list of Universal Resource Locator (URL) names of the sub-files. It is understood by one of skill in the art that it is possible to provide the configuration information through other techniques. For instance, it is possible to provide this information on a static or dynamically generated HTML web page. It is also possible to secure the configuration information by encrypting it.

The stitching program will use the configuration information to simultaneously request sub-file S1 from server 11.1.1.1, sub-file S2 from server 11.1.1.2, sub-file S3 from server 11.1.1.3, and sub-file S4 from server 11.1.1.4, as illustrated in Figure 8. Although this example shows one sub-file being transmitted from each server, it is possible for multiple sub-files to originate from a single server, with each sub-file transmitted on a separate TCP/IP connection.

As shown in Figure 9, sub-file S1 contains segments 1, 5, 9, and 13 of the original file. The chunk 1 data will be transmitted first, then chunk 5, and so on. Sub-file S2 contains segments 2, 6, 10, and 14 of the original file. The chunk 2 data will be transmitted first, then chunk 6, and so on. The segments of sub-file S1 will be transmitted separately and in parallel with the segments of sub-file S2, because each

sub-file would be transmitted over separate TCP/IP connections or via separate communication links.

In one embodiment, the sub-files will be transmitted at different rates from the different servers. In alternate embodiments, the sub-files may be transmitted at the same rates from one or more servers. Transmission errors, or lost packets for instance, will cause the TCP/IP protocol to retransmit a data packet until the packet is received correctly. This retransmission, among other things, will affect the transmission rate. Some paths may have more errors and lost packets than other paths. Some paths may also have more router hops and delays than other paths. This means that the reassembly of the segments into the original order of segments cannot be performed until the next required chunk has been fully and correctly received. For instance, once chunk 1 has been fully received and sent on to the media player, chunk 2 cannot be sent to the media player until it has been fully and correctly received, even if segments 3, 4, and 5 have been already received.

Once the media player has started, it must have data delivered at the proper rate, or the sound or video quality will suffer. To help guarantee that the data will be received in time, the stitching program measures the speed of each connection as data is received. If data is received at too slow of a rate on a connection, the slow connection can be terminated, and a new connection can be started to the fastest server that contains a copy of the required sub-file.

This server switching process may be implemented by providing multiple URLs in the configuration file. Each URL refers to a server that can provide the needed sub-file. An example of a URL list is shown in Figure 10A. The list has

three server choices for each sub-file. The initial server assignment for each sub-file is made from the first choice of possible assignments. The media transfers start out using these initial server assignments, as shown in Figure 10B. Sub-file S1.mpg is transferred from the server with the IP address of 11.1.1.1 (an exemplary IP address) using the HTTP protocol, sub-file S2.mpg is transferred from the server with the IP address of 11.1.1.2, sub-file S3.mpg is transferred from the server with the IP address of 11.1.1.3, and sub-file S4.mpg is transferred from the server with the IP address of 11.1.1.4.

At periodic intervals, the transfer rate from each server is measured and compared to a threshold for switching. If a sub-file's transfer rate is below the threshold, a new connection is established with the fastest server available for the sub-file. File offset information may be passed to the fast server about how much data had been received by the slow server, so that the fast server would not have to transmit duplicate information already received from the slow server. The connection to the slow server could be terminated once the new connection to the fast server has been established.

Among other options, the transfer rate switching threshold may be fixed, may be passed as a parameter in the configuration file, or may be made a function of other variables, such as the projected data buffer empty time. For instance, the transfer rate switching threshold may be increased if a large amount of data is buffered before it is needed, or the threshold may be decreased if only a small amount of data is buffered and the media player may soon be running out of input data.

Figure 10C illustrates one embodiment of how the server connections could be re-assigned, assuming that the connection from server 11.1.1.2 for sub-file S2.mpg and the connection from server 11.1.1.3 for sub-file S3.mpg were slow and below the transfer rate switching threshold. For sub-file S2.mpg, the server 11.1.1.1 is chosen for the new connection because that server is already providing data for sub-file A1.mpg at a rate above the transfer rate switching threshold. For sub-file S3.mpg, the server 11.1.1.4 is chosen for the new connection because that server is already providing data for sub-file S1.mpg at a rate above the transfer rate switching threshold.

This procedure allows data to be transferred from multiple servers simultaneously for higher transfer rates. The data would be recombined, or stitched back together, for the media player. Each path's data transfer rate may be periodically measured and the slow paths or servers may be switched out in favor of faster paths or servers. This switching may occur at the very start of the transfer, so that it would not be necessary to keep a continuously updated map of the Internet's congestion. The switching may also occur periodically during the media playback, so that unexpected slowdowns in data transfer from one or more servers or paths could be countered. Only a small amount of buffering will be needed to make the switching process transparent to the end user.

Figure 11 is a block diagram illustrating one example of a networked system 1100 that implements one embodiment of the invention. Client system 1100 running client-side software, is connected to one or more servers over network 30. Moreover, in this embodiment Servers A, B, C and D are running server-side software, and are

connected to network 30. However, it should be appreciated that client system 1110 and Servers A-D may be connected using any known communication link.

In one embodiment, the client-side software of client system 1110 includes stitching software 1120, which may be used to create reassembled file 1130. Thereafter, reassembled file 1130 may be provided to media presentation software 1140 for display. In one embodiment, stitching software 1120 is used to reassemble file segments which have been transferred from one or more of the Servers A-D. While any known means of reassembling file segments may be used, the method described in the PCT/US___/_____ patent application may be used, according to one embodiment.

Continuing to refer to Figure 11, sub-files S1 and S2 from Figure 4 have been loaded onto Server A. Similarly, sub-files S1, S3 and S4 have been loaded onto Server B, sub-files S1 and S3 have been loaded onto Server C and sub-files S2 and S4 have been loaded onto Server D. Thus, in this embodiment the original file 400 (see Figure 4) has been split into four sub-files S1 410, S2 420, S3 430 and S4 440, with each of the four sub-files S1-S4 being available from more than one server. As described earlier, this arrangement improves the speed at which a user will be able to download a given file.

The sub-files S1, S2, S3 and S4 may be accessed via network 30 upon demand. Upon retrieval, the files may be recombined using stitching software 1120 into reassembled file 1130. Thereafter, reassembled file 1130 may be provided to media presentation software 1140 for display.

While the preceding description has been directed to particular embodiments, it is understood that those skilled in the art may conceive modifications and/or variations to the specific embodiments described herein. Any such modifications or variations which fall within the purview of this description are intended to be included therein as well. It is understood that the description herein is intended to be illustrative only and is not intended to limit the scope of the invention.

CLAIMS

What is claimed is:

1. A method for receiving content over a network, where said content has been split into a plurality of segments and organized into a plurality of sub-files, the method comprising:

receiving sub-file data over the network, said sub-file data to identify a location of each of said plurality of sub-files;

requesting, using said sub-file data, a first sub-file and a second sub-file of said plurality of sub-files;

receiving over the network, in response to said request, said first sub-file and said second sub-file; and,

combining said first sub-file and second sub-file into a reassembled file.

2. The method of claim 1 wherein said receiving sub-file data over the network comprises receiving sub-file data over the network, where said sub-file data is to identify the location of each of said plurality of sub-files, and further to identify which of said plurality of segments are in each of said plurality of sub-files.

3. The method of claim 2, wherein said sub-file data is to further identify an order in which said plurality of segments are organized into said plurality of sub-files.

4. The method of claim 3, wherein said order is characterized by organizing a first segment into the first sub-file, a second segment into the second sub-file and a third segment into the first sub-file.

5. The method of claim 3, wherein said plurality of sub-files are further organized into a plurality of sub-file groups, and said order is characterized by organizing a first sequence of segments across a first sub-file group, followed by organizing a second sequence of segments across said first sub-file group.

6. The method of claim 1, wherein said receiving sub-file data is in response to a user request for said content.

7. The method of claim 1, wherein said receiving sub-file data over the network comprises receiving sub-file data over the network, said sub-file data to include a Universal Resource Locator (URL) for each of said plurality of sub-files.

8. The method of claim 1, wherein said requesting, using said sub-file data, the first sub-file and the second sub-file comprises requesting, using said sub-file data, the first sub-file from a first server and the second sub-file from a second server.

9. The method of claim 1, wherein receiving, in response to said request, said first sub-file and said second sub-file comprises receiving over the network, in response to said request, said first sub-file from a network server, and receiving over

the network, in response to said request, said second sub-file from said network server.

10. The method of claim 1, wherein receiving, in response to said request, said first sub-file and said second sub-file comprises receiving over the network, in response to said request, said first sub-file from a first server and said second sub-file from a second server.

11. The method of claim 1, wherein combining said first sub-file and second sub-file comprises combining said first sub-file and second sub-file into a reassembled file using the sub-file data.

12. The method of claim 11, further comprising presenting said reassembled file on a user computer using a media player.

13. An apparatus for receiving content over a network, where said content has been split into a plurality of segments and organized into a plurality of sub-files, the apparatus comprising:

a processor;

a communications port, coupled to the processor, said communications port to be used to connect said apparatus to a network; and

a memory, coupled to the processor, said memory to include instruction sequences to cause said processor to,

receive sub-file data over the network, said sub-file data to identify a location each of said plurality of sub-files,
request, using said sub-file data, a first sub-file and a second sub-file of said plurality of sub-files,
receive over the network, in response to said request, said first sub-file and said second sub-file, and,
combine said first sub-file and second sub-file into a reassembled file.

14. The apparatus of claim 13, wherein said sub-file data further identifies which of said plurality of segments are in each of said plurality of sub-files.

15. The apparatus of claim 14, wherein said sub-file data is to further identify an order in which said plurality of segments are organized into said plurality of sub-files.

16. The apparatus of claim 15, wherein said order is characterized by organizing a first segment into the first sub-file, a second segment into the second sub-file and a third segment into the first sub-file.

17. The apparatus of claim 15, wherein said plurality of sub-files are further organized into a plurality of sub-file groups, and said order is characterized by organizing a first sequence of segments across a first sub-file group, followed by organizing a second sequence of segments across said first sub-file group.

18. The apparatus of claim 13, wherein said sub-file data is to be received in response to a user request for said content.

19. The apparatus of claim 13, wherein said sub-file data is to include a Universal Resource Locator (URL) for each of said plurality of sub-files.

20. The apparatus of claim 13, wherein said memory further includes instruction sequences to cause said processor to request, using said sub-file data, the first sub-file from a first server and the second sub-file from a second server.

21. The apparatus of claim 13, wherein said memory further includes instruction sequences to cause said processor to receive over the network, in response to said request, said first sub-file from a network server, and to receive over the network, in response to said request, said second sub-file from said network server.

22. The apparatus of claim 13, wherein said memory further includes instruction sequences to cause said processor to receive, in response to said request, said first sub-file from a first server and said second sub-file from a second server.

23. The apparatus of claim 13, wherein said memory further includes instruction sequences to cause said processor to combine said first sub-file and second sub-file into a reassembled file using the sub-file data.

24. The apparatus of claim 23, wherein said memory further includes instruction sequences to cause said processor to present said reassembled file using a media player.

25. A computer program product, comprising:
a computer usable medium having computer readable program code embodied therein to provide content over a network where the content has been split into a plurality of segments and organized into a plurality of sub-files, the computer readable program code in said computer readable program product comprising:

first computer readable program code to receive sub-file data over the network, said sub-file data to identify a location of each of said plurality of sub-files;

second computer readable program code to issue a request, using said sub-file data, for a first sub-file and a second sub-file of said plurality of sub-files;

third computer readable program code to receive over the network, in response to said request, said first sub-file and said second sub-file; and,

fourth computer readable program code to combine said first sub-file and second sub-file into a reassembled file.

26. The computer program product of claim 25, wherein the first computer readable program code comprises first computer readable program code to receive sub-file data over the network, where said sub-file data is to identify the location of

each of said plurality of sub-files, and further to identify which of said plurality of segments are in each of said plurality of sub-files.

27. The computer program product of claim 26, wherein the first computer readable program code comprises first computer readable program code to receive sub-file data over the network, where said sub-file data is to identify the location of each of said plurality of sub-files, identify which of said plurality of segments are in each of said plurality of sub-files, and to identify an order in which said plurality of segments are organized into said plurality of sub-files.

28. The apparatus of claim 27, wherein said order is characterized by organizing a first segment into the first sub-file, a second segment into the second sub-file and a third segment into the first sub-file.

29. The apparatus of claim 27, wherein said plurality of sub-files are further organized into a plurality of sub-file groups, and said order is characterized by organizing a first sequence of segments across a first sub-file group, followed by organizing a second sequence of segments across said first sub-file group.

30. The computer program product of claim 25, wherein the first computer readable program code comprises first computer readable program code to receive sub-file data over the network in response to a user request for said content, said sub-file data to identify a location of each of said plurality of sub-files.

31. The computer program product of claim 25, wherein the first computer readable program code comprises first computer readable program code to receive sub-file data over the network, said sub-file data to include a Universal Resource Locator (URL) for each of said plurality of sub-files.

32. The computer program product of claim 25, wherein the second computer readable program code comprises second computer readable program code to issue the request, using said sub-file data, for the first sub-file from a first server and the second sub-file from a second server.

33. The computer program product of claim 25, wherein the third computer readable program code comprises third computer readable program code to receive over the network, in response to said request, said first sub-file from a network server, and to receive over the network, in response to said request, said second sub-file from said network server.

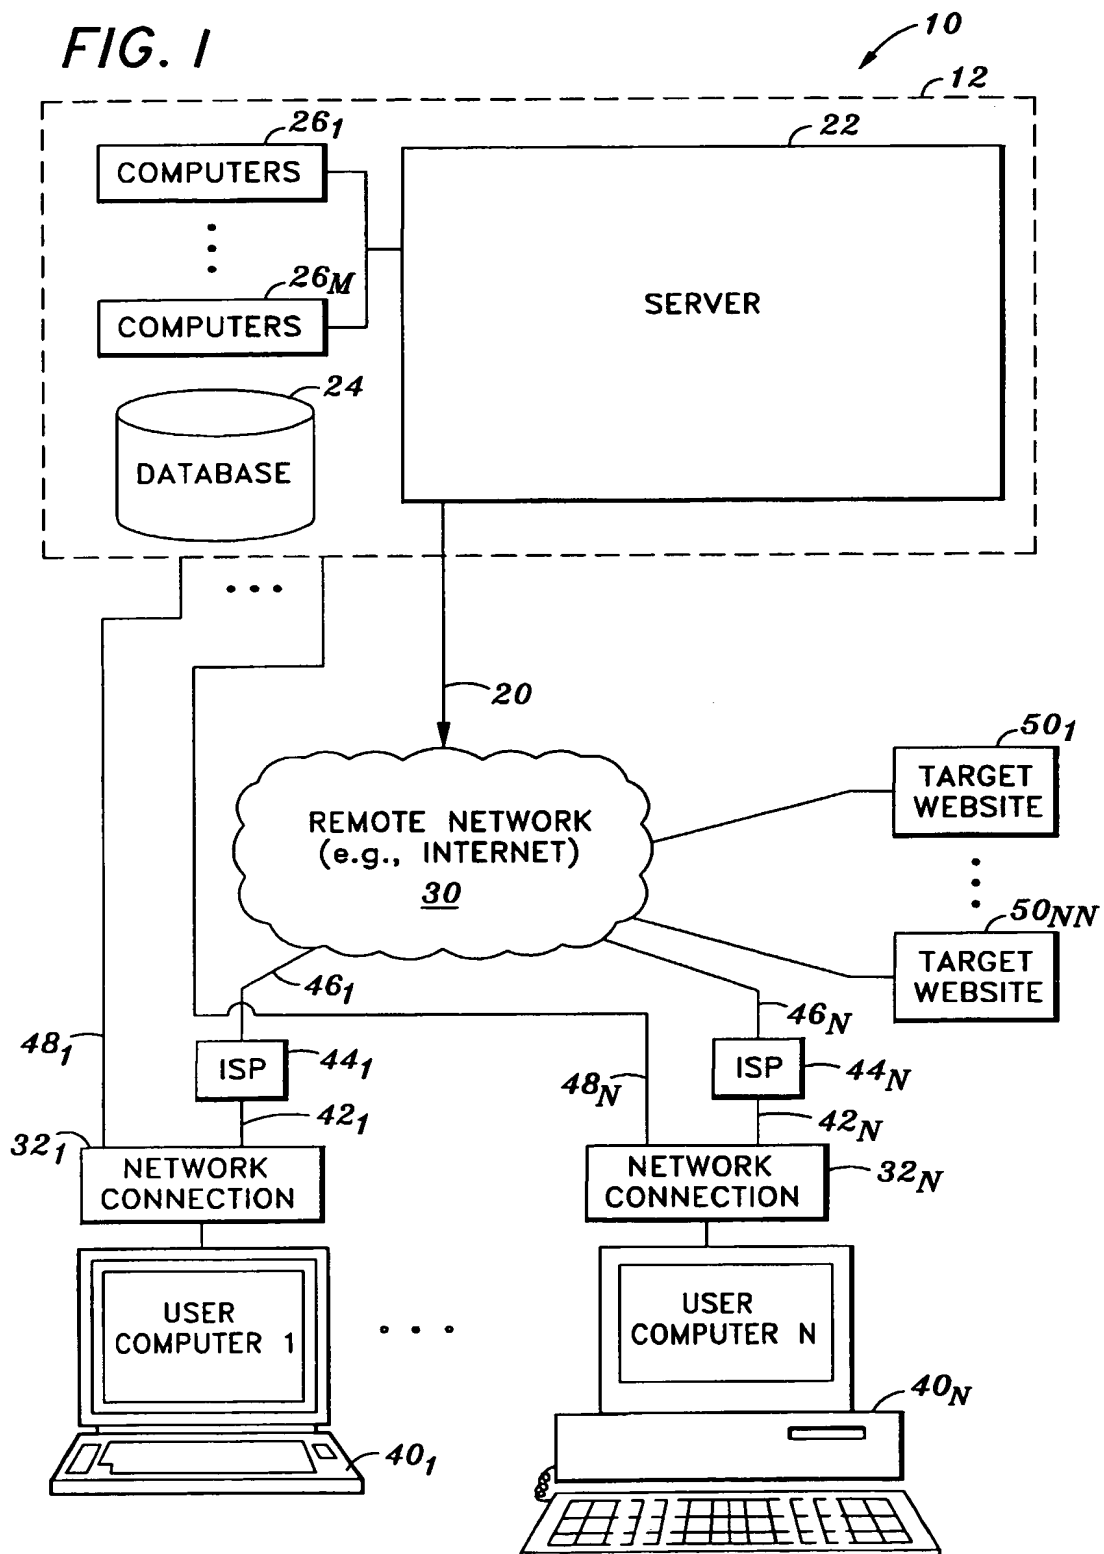
34. The computer program product of claim 25, wherein the third computer readable program code comprises third computer readable program code to receive over the network, in response to said request, said first sub-file from a first server and said second sub-file from a second server.

35. The computer program product of claim 25, wherein the fourth computer readable program code comprises fourth computer readable program code

to combine said first sub-file and second sub-file into a reassembled file using the sub-file data.

36. The computer program product of claim 35, further comprising fifth computer readable program code to present said reassembled file on a user computer using a media player.

FIG. 1






FIG. 3

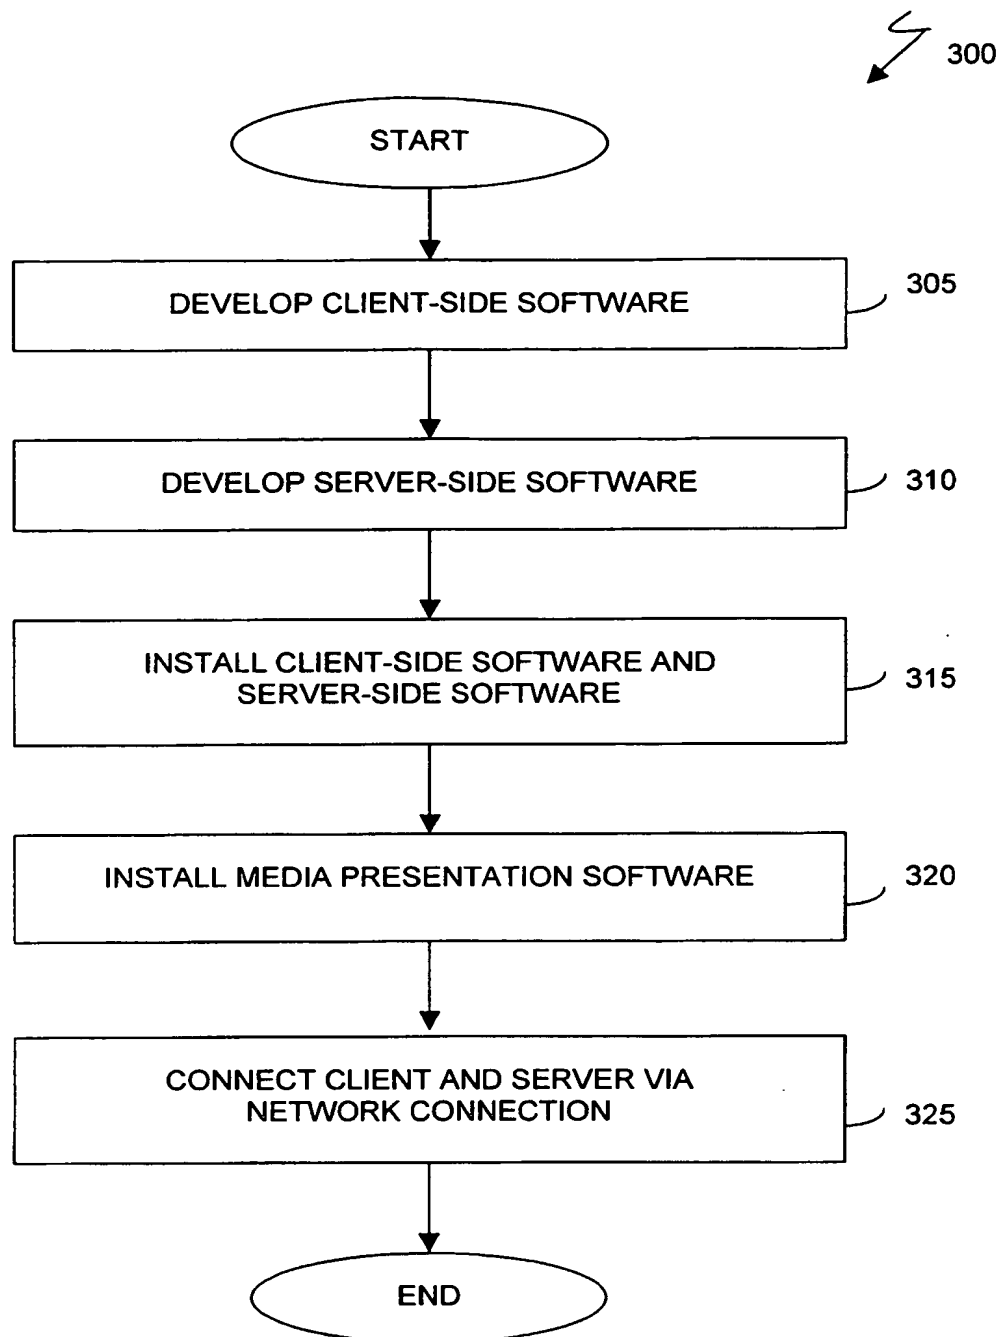


FIG 4

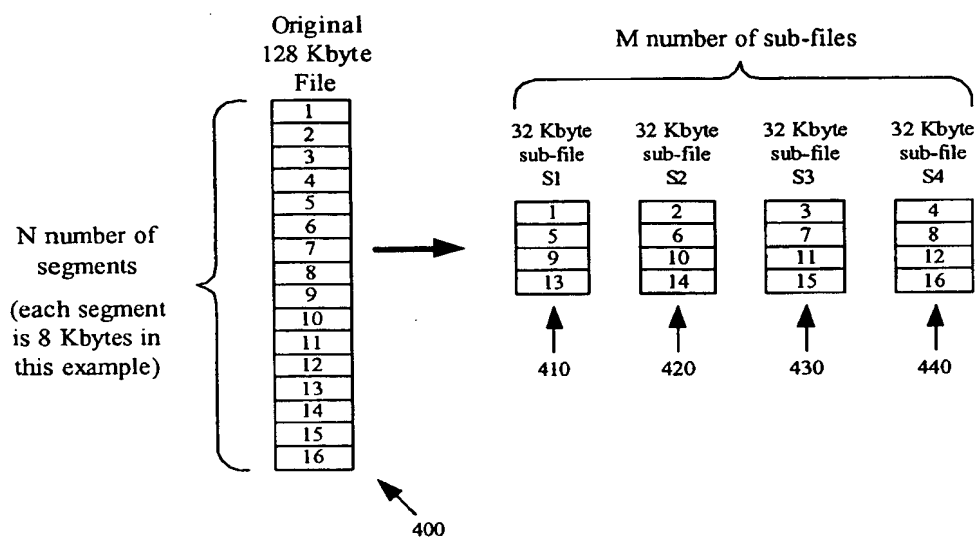


FIG 5




FIG 6

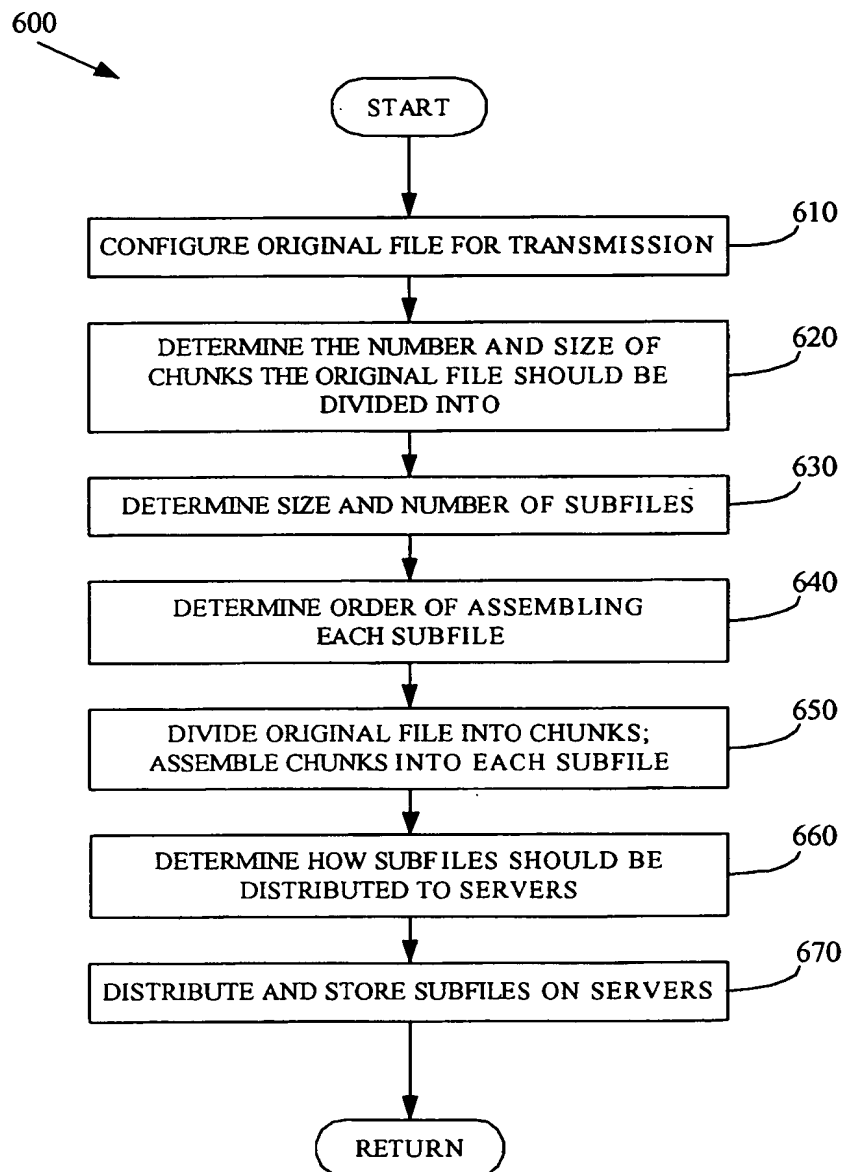


FIG 7




FIG 8

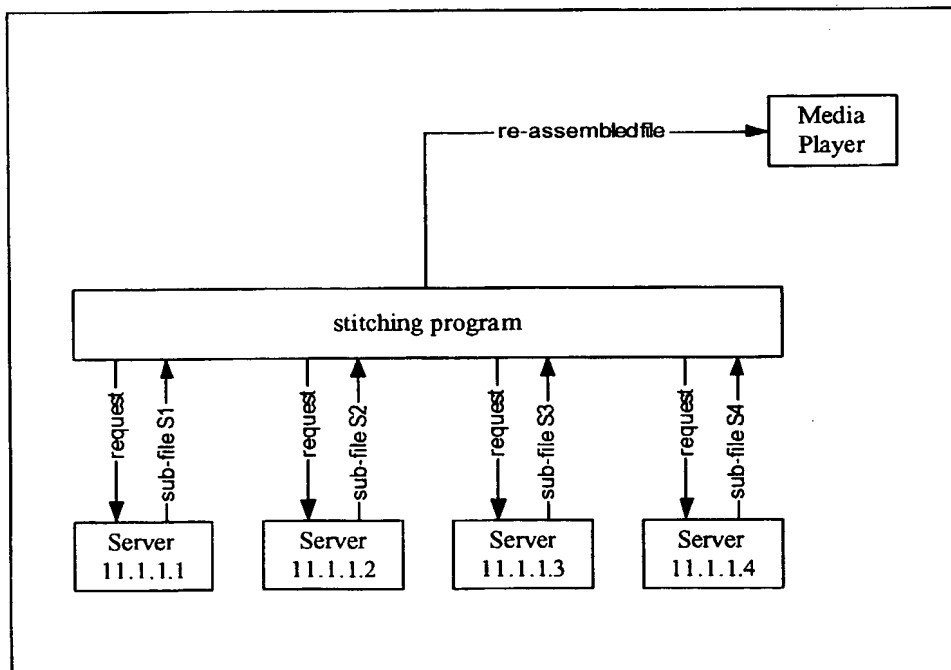


FIG 9




FIG 10

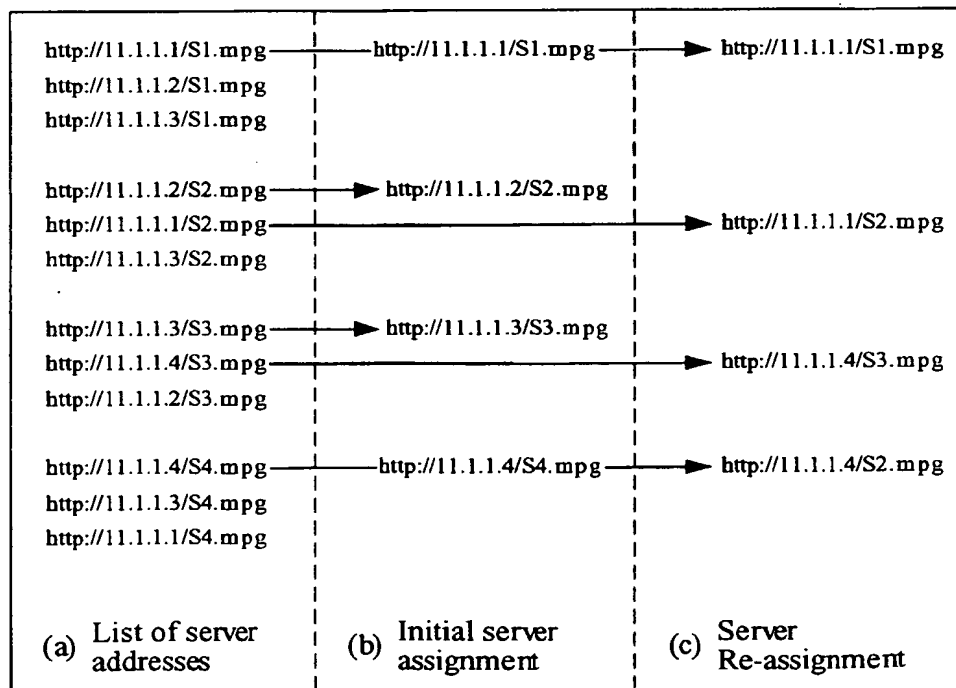
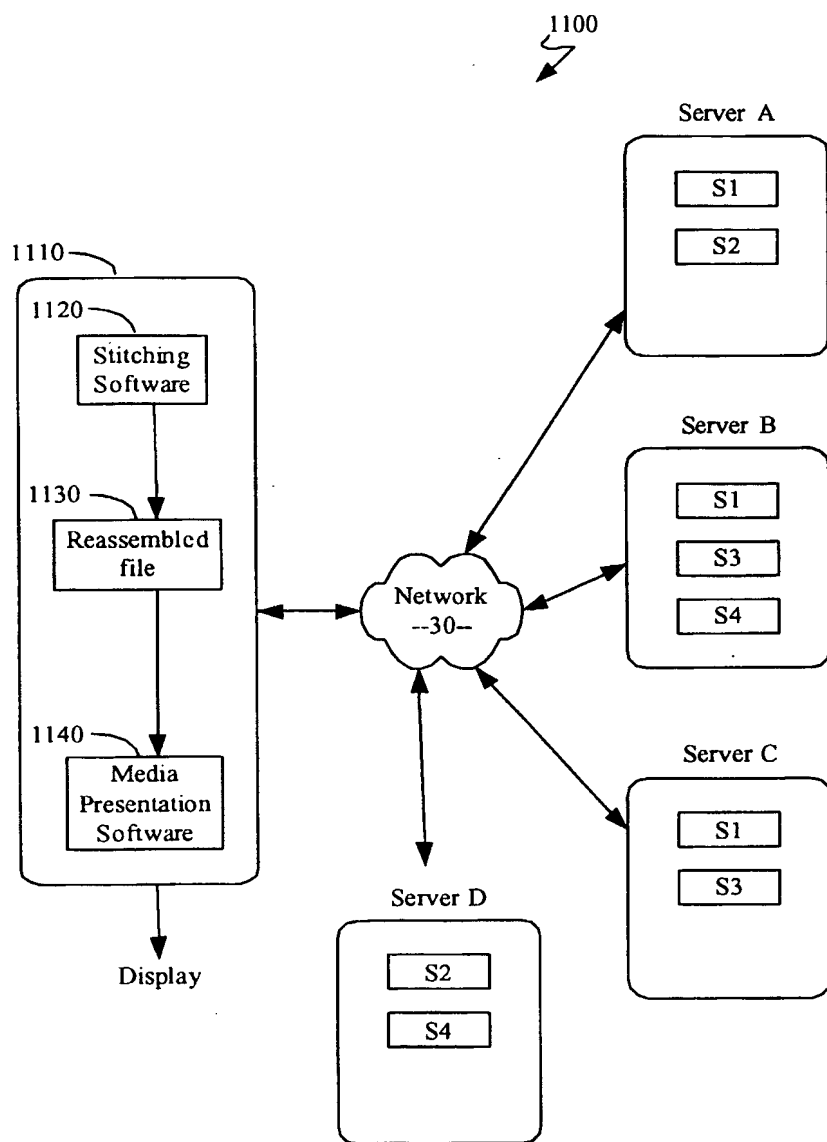


FIG 11



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US02/07659

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 15/16

US CL : 709/219

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 709/219, 725/89, 90

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EAST IEEE ONLINE

search terms: video streaming, segmented, transmission, instructions

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X ----- Y	WO 98/33320 A1 (EVERETT et al.) 30 JULY 1998 PAGES 4-6	1-6, 11-17, 21-27, 31 ----- 7-10, 18-20, 28-30
X ----- Y	WO 97/28620 A1 (SOLOMAKI) 7 AUGUST 1997 PAGES 8-16, FIGURE 9	1-6, 11-17, 21-27, 31 ----- 7-10, 18-20, 28-30
X ----- Y	US 6,085,251 A (FABOZZI, II) 4 JULY 2000,	1-6, 11-17, 21-27, 31 ----- 7-10, 18-20, 28-30



Further documents are listed in the continuation of Box C.



See patent family annex.

<p>* Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p>		<p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art</p> <p>"Z" document member of the same patent family</p>
--	--	---

Date of the actual completion of the international search

04 JUNE 2002

Date of mailing of the international search report

26 JUN 2002

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

LARRY DONAGUE

Telephone No. (703) 305-9675

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US02/07659

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 6,170,060 B1 (MOTT et al.) 2 JANUARY 2001, col. 15, line 57 - col. 16, line 20	7-10, 18-20, 28-30